arXiv:1609.02901v1 [math.PR] 9 Sep 2016

RAPID MIXING OF GEODESIC WALKS ON MANIFOLDS WITH POSITIVE CURVATURE

OREN MANGOUBI* AND AARON SMITH†

ABSTRACT. We introduce a Markov chain for sampling from the uniform distribution on a Riemannian manifold $\mathcal{M}$, which we call the *geodesic walk*. We prove that the mixing time of this walk on any manifold with positive sectional curvature $C_x(u,v)$ bounded both above and below by $0 < \mathfrak{m}_2 \leq C_x(u,v) \leq \mathfrak{M}_2 < \infty$ is $\mathcal{O}^*\left(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}\right)$. In particular, this bound on the mixing time does not depend explicitly on the dimension of the manifold. In the special case that $\mathcal{M}$ is the boundary of a convex body, we give an explicit and computationally tractable algorithm for approximating the exact geodesic walk. As a consequence, we obtain an algorithm for sampling uniformly from the surface of a convex body that has running time bounded solely in terms of the curvature of the body.

1. INTRODUCTION

Sampling from manifolds has many applications to areas such as statistics [6], computer graphics [15], and systems biology [18]. In this paper we study a simple Markov chain $\{X_i\}_{i\in\mathbb{N}}$, which we call the *geodesic walk*, that can sample efficiently from the uniform distribution of a general Riemannian manifold $\mathcal{M}$. The geodesic walk, defined precisely in Algorithm 1 below, evolves by selecting at each time step a random tangent vector U_i in the tangent space $\mathcal{T}_{X_i}$ of $\mathcal{M}$ at X_i, and then following the associated geodesic for some period of time.

1.1. Main Results.

The geodesic walk is a natural Markov chain on a manifold, and it is somewhat similar to the well-studied "ball walk" on a manifold (see *e.g.* [12]). It is well-known (see Example 7 of [14]) that the ball walk mixes rapidly on manifolds with positive curvature. One main result of this paper, Theorem 5, is that the geodesic walk also mixes rapidly on a manifold with positive curvature.

*omangoubi@gmail.com, Department of Mathematics and Statistics, University of Ottawa, 585 King Edward Avenue, Ottawa ON K1N 7N5, Canada.

†smith.aaron.matthew@gmail.com, Department of Mathematics and Statistics, University of Ottawa, 585 King Edward Avenue, Ottawa ON K1N 7N5, Canada.

Although the ball walk is easy to define and mixes rapidly, it is very difficult to implement on a computer and does not give rise to practical sampling algorithms (see Example 3C: *How not to sample* of [6]). In contrast, we show that the geodesic walk is a useful general tool for sampling. Our second main result (Theorems 1 and 4) is that, unlike the ball walk, the geodesic walk has a unique stationary distribution that is uniform with respect to the volume measure on the manifold.

The geodesic walk is also much easier to implement. Our third main result, Theorem 6, shows that any reasonably good algorithm for approximating geodesics can be used to simulate an "approximate" geodesic walk that mixes quickly and has near-uniform stationary measure.

We believe our $\mathcal{O}^*\left(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}\right)$ bound on the mixing time of the geodesic walk is the first dimensionless mixing time bound for an implementable Markov chain with uniform stationary distribution on bounded-positive curvature manifolds in general, and convex body boundaries in particular. One reason our bound does not depend on the dimension is that, in contrast to previously-proposed algorithms, the geodesic walk can take long steps whose length is independent of the dimension. For comparison, the best existing bound that we know of is given for the "billiards walk" studied in [7], which is shown to have mixing time of $\mathcal{O}^*\left(d^2 \cdot \frac{\mathfrak{M}_2}{\mathfrak{m}_2}\right)$.[1]

In both statistics and computer science, one of the most important and well-studied special cases occurs when $\mathcal{M}$ is the boundary $\partial\mathcal{K}$ of a convex body $\mathcal{K} \subset \mathbb{R}^{d+1}$. There are many Markov chains for sampling from the uniform distribution on the boundary $\partial\mathcal{K}$ of a set under various assumptions (see *e.g.* [2], [5], [7]), all of which involve a walk that moves inside of $\mathcal{K}$ and then reflects off of $\partial\mathcal{K}$. In Section 8, we show that our geodesic walk can be well-approximated by such a "reflecting walk," if the reflecting walk is allowed to have long-term momentum. Our final main result, Theorem 7, shows that this approximate version of the geodesic walk gives us an algorithm for sampling uniformly from the boundary $\partial\mathcal{K}$ of a convex set $\mathcal{K}$ with arbitrarily small error in roughly $\mathcal{O}^*\left((\frac{\mathfrak{M}_2}{\mathfrak{m}_2})^3\right)$ reflections, provided that $\mathcal{M}$ has inner and outer radii of curvature bounded below and above by $\frac{1}{\sqrt{\mathfrak{M}_2}}$ and $\frac{1}{\sqrt{\mathfrak{m}_2}}$,

[1]Note that the bound in [7] is better for chains with a bound on the diameter of $\mathcal{K}$ that is much smaller than $\frac{2}{\sqrt{\mathfrak{m}_2}}$, while our bound is better in higher dimensions and when the bound on the inner radius of curvature of $\mathcal{K}$ is much smaller than $\frac{1}{\sqrt{\mathfrak{M}_2}}$. Also, a practical implementation of the geodesic walk requires additional computational costs that depend on the method used to approximate geodesics.

respectively. We believe the $\mathcal{O}^*\left((\frac{\mathfrak{M}_2}{\mathfrak{m}_2})^3\right)$ bound for the implementation of the geodesic walk is often very pessimistic and can be greatly improved in some important special cases (See Remark 2 in Section 8.2). Furthermore, in situations where one has access to better geodesic approximations, such as when symplectic integrators or exact geodesic integrators are available, our results imply computational costs that are much closer to the mixing time $\mathcal{O}^*\left(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}\right)$ of the true geodesic walk.

1.2. **Proof Techniques.** The main component of our proof is a coupling argument, where we couple the initial velocity for each geodesic trajectory of two geodesic walks so that the initial velocity of one chain is the parallel transport of the other chain's initial velocity along a shortest geodesic connecting the two chains (Section 5). We then use comparison theorems from differential geometry [1, 4, 17] to show that our assumptions of positive curvature bounds imply that the distance between the two chains contracts over each step in the Markov chain (Section 6). We then use arguments similar to those used in [14] to bound the Wasserstein transportation distance mixing time of the geodesic walk (Section 7). Finally, we show that one can computationally approximate geodesic trajectories with arbitrary accuracy in a dimension-independent number of steps provided one has access to an appropriate oracle, allowing one to generate samples from a stationary distribution that is arbitrarily close to uniform (in Wasserstein transportation distance) in a dimension-independent number of oracle calls. We also show how to construct such an oracle in the special case where $\mathcal{M}$ is the boundary of a convex body (Section 8).

Finally, we mention that there is also a mixing time result for a related (but different) geodesic walk [19] in a very different setting where one wishes to sample from the boundary of a convex polytope. However, the mixing time bound for the geodesic walk in [19] is not dimension-independent. Another difference is that the authors of [19] state that their proof follows the conductance-based approach that is popular for analyzing geometric random walks [20], while our proof uses a probabilistic coupling argument as in [14].

1.3. **Motivations.** The geodesic walk is a natural choice of Markov chain to generate uniform samples from the volume measure on a manifold. We highlight three convenient properties of the geodesic walk:

(1) The geodesic walk samples directly from the uniform distribution, due to the symplectic properties of geodesic flow. In contrast, most other simple walks on manifolds, such as the ball

walk, must be adjusted by the addition of a Metropolis-Hastings correction in order to sample from the uniform distribution. This correction slows down mixing, and can be computationally intractable in high dimensions.

(2) The geodesic walk takes very long steps, allowing it to mix rapidly. If the manifold has bounded positive curvature, we can prove rapid mixing bounds.

(3) The geodesic walk is easy to implement in a computational setting, by computing an approximation to its geodesic paths.

1.4. **List of Sections.** The rest of the paper is arranged as follows:

- In Section 2 we go over the assumptions we make about $\mathcal{M}$ and Riemannian geometry preliminaries.
- In Section 3 we define the geodesic walk on the manifold $\mathcal{M}$.
- In Section 4, we prove that the stationary distribution of the geodesic walk is uniform on $\mathcal{M}$.
- In Section 5 we define a coupling of two copies of the geodesic walk.
- In Section 6 we prove the contraction bound described above.
- In Section 7 we use this contraction bound to prove a bound on the mixing time.
- In Section 8 we show that one can computationally approximate geodesic trajectories to sample the uniform distribution on $\mathcal{M}$ to arbitrary accuracy in a dimension-independent number of steps if one has access to an appropriate oracle (Section 8.1). We construct such an oracle explicitly in the special case where $\mathcal{M}$ is the boundary of a convex body (Section 8.2).

2. Assumptions and Riemannian geometry preliminaries

In this section we go over the assumptions about our manifold $\mathcal{M}$, and Riemannian geometry preliminaries, and introduce some definitions we will use in the rest of the paper.

- **Assumptions about our manifold $\mathcal{M}$:** Throughout this paper we will assume that,
  (1) $(\mathcal{M}, g)$ is a closed, connected, second-order differentiable Riemannian manifold, with associated inner product $g \equiv g_x$ on the tangent space $\mathcal{T}_x \equiv \mathcal{T}_x\mathcal{M}$ at $x \in \mathcal{M}$.
  (2) $(\mathcal{M}, g)$ has bounded positive curvature. That is, there exist constants $0 < \mathfrak{m}_2 \leq \mathfrak{M}_2 < \infty$ so that, for all $x \in \mathcal{M}$ and $u, v \in \mathcal{T}_x$,

$$\mathfrak{m}_2 \leq C_x(u, v) \leq \mathfrak{M}_2, \tag{1}$$

where $C_x(u,v)$ is the sectional curvature of $\mathcal{M}$ at x in the directions $\frac{u}{\|u\|}$ and $\frac{v}{\|v\|}$.

- **Levi-Civita connection:** The fundamental theorem of Riemannian geometry guarantees the existence and uniqueness of a torsion-free affine connection ∇ on $(\mathcal{M}, g)$ that induces an isometry of tangent spaces via parallel transport (see Theorem 6.8 of [9]); this connection is called the *Levi-Civita connection*.
- **Paths and Metrics:** Recall that the length of a segment of a smooth path $h : \mathbb{R}^+ \mapsto \mathcal{M}$ is given by

$$\text{length}(h[a,b]) = \int_a^b \sqrt{g_{h(t)}(h'(t), h'(t))}\mathrm{d}t. \tag{2}$$

This gives rise to the metric

$$\text{dist}(x,y) = \inf_{h\,:\,h(0)=x,\,h(1)=y} \text{length}(h[0,1]) \tag{3}$$

on $\mathcal{M}$.

- **Uniform Measures:** We recall two measures, one on the manifold $\mathcal{M}$ and the other on small subsets of the tangent spaces $\mathcal{T}_x$, that we will call the *uniform* measures on their associated spaces. First, recall that any Riemannian manifold $(\mathcal{M}, g)$ has an associated volume function λ. If $\lambda(\mathcal{M}) < \infty$, we denote by $\text{Unif}(\mathcal{M})$ the probability measure given by

$$\text{Unif}(\mathcal{M})[A] = \frac{\lambda(A)}{\lambda(\mathcal{M})} \tag{4}$$

for $A \subset \mathcal{M}$ measurable. We refer to this as the *uniform measure on* $\mathcal{M}$.

Recall that any choice of the basis $\mathcal{B}$ of the tangent space $\mathcal{T}_x$ of a d-dimensional manifold $\mathcal{M}$ at a point $x \in \mathcal{M}$ gives a natural isometry $\zeta_x = \zeta_{x,\mathcal{B}} : \mathcal{T}_x \mapsto \mathbb{R}^d$. Furthermore, the pull-back of the Lebesgue measure from the unit sphere $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\} \subset \mathbb{R}^d$ to the set $\mathbb{S}(\mathcal{T}_x) = \zeta_x^{-1}(\mathbb{S})$ given by ζ_x does not depend on the choice of basis $\mathcal{B}$. We denote by $\text{Unif}(\mathbb{S}(\mathcal{T}_x))$ this unique pull-back of the Lebesgue measure by ζ_x, and refer to this measure as the *uniform measure on the unit sphere in* $\mathcal{T}_x$.

Throughout the paper, if $S_1 \subset S_2$ and we have defined a uniform measure on S_2, we define the uniform measure on S_1 by

$$\text{Unif}(S_1)[A] = \frac{\text{Unif}(S_2)[A]}{\text{Unif}(S_2)[S_1]}.$$

- **Geodesics:** Define the phase space

$$\mathcal{M}^\circ := \{(x, v) : x \in \mathcal{M}, v \in \mathbb{S}(\mathcal{T}_x)\}. \tag{5}$$

Associated to every element $(x, v) \in \mathcal{M}^\circ$, there is a special path $\gamma_{(x,v)} : \mathbb{R}^+ \to \mathcal{M}$ with $\gamma_{(x,v)}(0) = x$ and $\gamma'_{(x,v)}(0) = v$ that is called the *geodesic*. Roughly speaking, this is the path obtained by starting at point x and traveling with velocity v along the manifold. We will use the following properties of the geodesic [16]:

(1) For all $x \in \mathcal{M}$ and $v \in \mathcal{T}_x$, there exists $\epsilon = \epsilon(x, v) > 0$ so that

$$\text{length}(\gamma_{(x,v)}[0, t]) = \text{dist}(\gamma_{(x,v)}(0), \gamma_{(x,v)}(t))$$

for all $0 < t < \epsilon$. (See Section 5 and Exercise 5.9.34 of [16])

(2) Associated to every smooth map $\tau : \mathbb{R}^+ \mapsto \mathcal{M}^\circ$, there is an object called the *Jacobi field*

$$J_\tau(t) := \frac{\partial \gamma_\tau}{\partial \tau}, \tag{6}$$

that satisfies the Jacobi second-order ordinary differential equation:

$$\frac{\mathrm{D}^2}{\mathrm{d}t^2} J_\tau(t) + R(J(t), J'(t))J'(t) = 0, \tag{7}$$

where R denotes the Riemann curvature tensor of $\mathcal{M}$. (See Section 3.2.4 of [16])

- **Diameter:** We define $D := \text{Diam}(\mathcal{M})$ to be the diameter of $\mathcal{M}$. Since the sectional curvature is bounded below by $\mathfrak{m}_2$, the Myers-Bonnet theorem [13] implies that

$$D \leq \frac{\pi}{\sqrt{\mathfrak{m}_2}}. \tag{8}$$

- **A *unique* parallel transport:** Associated with any *connection* on a manifold and any finite-length path $h : [a, b] \mapsto \mathcal{M}$ there exists a map from $\mathcal{T}_{h(a)}$ to $\mathcal{T}_{h(b)}$, called the *parallel transport* map.

We now define the only parallel transport maps that we will use in this paper. The Hopf-Rinow theorem [10] implies that between any pair of points $x, y \in \mathcal{M}$ there exists at least one unit-speed geodesic $\omega(t) \equiv \omega(t; x, y)$, $t \in [0, \text{dist}(x, y)]$ with $\omega(0; x, y) = x$ and $\omega(\text{dist}(x, y); x, y) = y$ with $\omega(t; x, y) = \omega(\text{dist}(x, y) - t; y, x)$. We choose arbitrarily, via the axiom of choice, a particular family $\{\omega(\cdot; x, y)\}_{(x,y)\in\mathcal{M}\times\mathcal{M}}$ of such geodesic paths. Finally, we define $\phi(t; \cdot; x, y) : \mathcal{T}_x \mapsto \mathcal{T}_{\omega(t;x,y)}$ to

be the parallel transport map associated with the Levi-Civita connection and the path $\omega(\cdot; x, y)$. We also write $\phi(\cdot; x, y) \equiv \phi(\mathrm{dist}(x, y); \cdot; x, y)$ for shorthand.

Finally, fix an arbitrary point $x \in \mathcal{M}$ and an arbitrary basis $\{b_i\}$ of $\mathcal{T}_x$. Throughout the paper, we denote by $\{\zeta_y\}_{y\in\mathcal{M}}$ the maps $\zeta_{y,\{\phi(b_i;x,y)\}} : \mathcal{T}_y \mapsto \mathbb{R}^d$ associated with this basis and its parallel transports by ϕ.

- **Wasserstein distance and mixing time:** For any distribution μ, we write $X \sim \mu$ when the random variable X has distribution μ. For two distributions μ, ν on a common measure space $(\Omega, \mathcal{A})$, define $\Xi(\mu, \nu)$ to be the collection of all distributions ξ on $(\Omega^2, \mathcal{A}\times\mathcal{A})$ with marginal distributions $\xi(\cdot, \Omega) = \mu(\cdot)$, $\xi(\Omega, \cdot) = \nu(\cdot)$. The Wasserstein transportation distance W_d between two measures μ and ν on a common metric measure space (Ω, d) is given by

$$W_d(\mu, \nu) = \inf_{(X,Y)\,:\,(X,Y)\sim\Xi(\mu,\nu)} \mathbb{E}[d(X, Y)].$$

  Consider a transition kernel K on a metric measure space (Ω, d), with unique stationary distribution ξ. We define the Wasserstein mixing profile $\mathfrak{t}_{\mathrm{mix}} : [0, 1] \mapsto \mathbb{N}$ of K to be

$$\mathfrak{t}_{\mathrm{mix}}(\epsilon) = \inf\{\mathfrak{t} \geq 0 : \sup_{x\in\Omega} W_d(K^{\mathfrak{t}}(x, \cdot) - \xi(\cdot)) < \epsilon\}.$$

- In this paper we use the "big-$\mathcal{O}$" notation. Specifically, for any two functions $f : \mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R} \to \mathbb{R}$, we write "$f(z) = \mathcal{O}(g(z))$" or "$f(z) \leq \mathcal{O}(g(z))$" if there exists a constant $0 < C < \infty$ and some $Z > 0$ such that

$$f(z) \leq C \cdot g(z)$$

  for all $z > Z$. Furthermore, write "$f(z) = \mathcal{O}^*(g(z))$" or "$f(z) \leq \mathcal{O}^*(g(z))$" if there exist constants $0 < c', C' < \infty$ and some $Z' > 0$ such that

$$f(z) \leq C' \cdot g(z) \log(z)^{c'}$$

  for all $z > Z'$. I.e., the "big-$\mathcal{O}^*$" notation suppresses the logarithmic terms that we would otherwise need to take into account when using the "big-$\mathcal{O}$" notation.

3. The geodesic walk

The geodesic walk $\{X_i\}_{i\in\mathbb{N}}$ on a manifold (Figure 1), with a fixed geodesic step size T is defined precisely in Algorithm 1:

Algorithm 1 Geodesic Walk

parameters: Integration time T, manifold $(\mathcal{M}, g)$, number of steps N.
input: $X_1 \in \mathcal{M}$.
output: First N steps $\{X_i\}_{i=1}^N$ of the geodesic walk Markov chain on $\mathcal{M}$.

1: **for** $i = 1, 2, \ldots, N-1$ **do**
2: Sample $U_i \in \mathcal{T}_{X_i}$ from the uniform distribution $\mathrm{Unif}(\mathbb{S}(\mathcal{T}_{X_i}))$ on the unit sphere in $\mathcal{T}_{X_i}$.
3: Set $X_{i+1} = \gamma_{(X_i, U_i)}(T)$.
4: **end for**

That is, the point X_{i+1} is generated from X_i by running a geodesic trajectory with initial conditions (X_i, U_i) for a fixed time T.

We define the transition kernel $\mathbb{K}$ of the geodesic walk Markov chain $\{X_i\}_{i\in\mathbb{N}}$ by:

$$\mathbb{K}(Q, R) := \mathbb{P}[X_{i+1} \in R | X_i \sim \mathrm{Unif}(Q)] \tag{9}$$

for all measurable subsets $Q, R \subset \mathcal{M}$.

For the geodesic walk $x = X_1, X_2, \ldots$, we define $\Pi_i(\cdot) = \mathbb{K}^i(x, \cdot)$ to be the distribution of X_i, and Π to be the stationary distribution of the geodesic walk. In Section 4 we will prove that the uniform measure on $\mathcal{M}$ is a stationary distribution of the geodesic walk (later, we show that this is the only stationary measure).

We note that, in every step of Algorithm 1, a random variable U_i was constructed. It is straightforward to see that the sequence $\{(X_i, U_i)\}_{i\in\mathbb{N}}$ is a Markov chain on $\mathcal{M}^\circ$, which we will call the *phase-space* Markov chain.

Although we define the geodesic walk as an "algorithm", we will assume that we have an oracle for computing geodesic trajectories γ with perfect accuracy, which we refer to as an "idealized geodesic integrator" (We drop this assumption in Section 8, where we discuss a computational implementation of the geodesic walk).

4. The stationary distribution of the geodesic walk

The aim of this section is to prove that the transition kernel $\mathbb{K}$ defined in Equation 9 has uniform stationary distribution. To do so we will use the fact that the Liouville measure is invariant under geodesic flow [3, 11, 22].

Recall the definition of the phase space $\mathcal{M}^\circ$ in Equation 5. The Liouville measure $\mathcal{L}$ is defined to be the measure on $\mathcal{M}^\circ$ with density

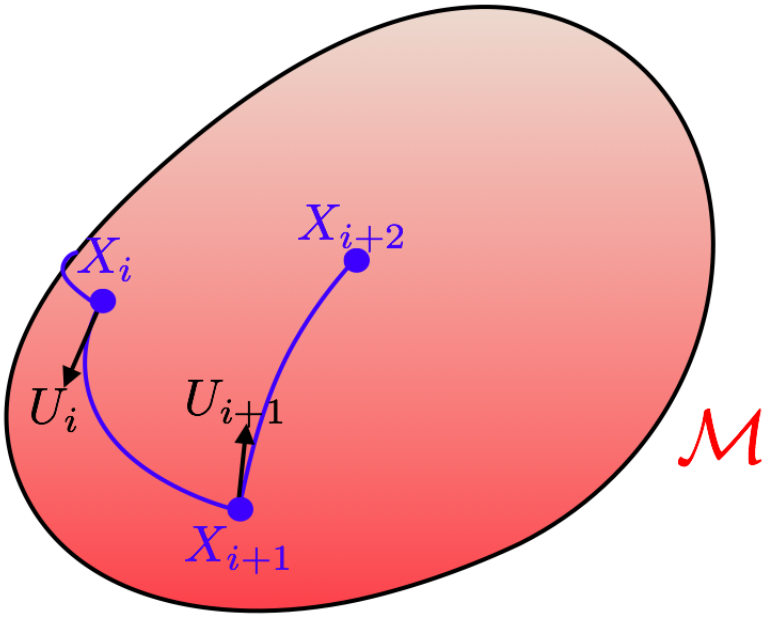

FIGURE 1. The geodesic walk $X_1, X_2, \ldots$ (blue). Given the current point X_i, the next point is generated by independently sampling a uniform random velocity U_i from the unit sphere on the tangent space at X_i, and running a geodesic trajectory with initial conditions (X_i, U_i) for a fixed time T.

given by the product of the volume form $\mathrm{d}\lambda(x)$ on the manifold and the volume form $\mathrm{d}\mu(v)$ on the unit sphere in the tangent space of the manifold at the point x (i.e., the unit-speed velocities):

$$\mathcal{L}(\Omega) := \int_{\Omega} \mathrm{d}\mu(v)\mathrm{d}\lambda(x). \tag{10}$$

Consider a geodesic trajectory on $\mathcal{M}$ with initial position x and initial velocity v. We define $\gamma_{(x,v)}(t)$ and $\varphi_{(x,v)}(t) := \gamma'_{(x,v)}(t)$ to be, respectively, the position and velocity of this trajectory at time $t \in \mathbb{R}^+$. Finally, we define

$$\psi_{(x,v)}(t) := (\gamma_{(x,v)}(t), \varphi_{(x,v)}(t))$$

to be the location of our geodesic trajectory in the phase space at time $t \in \mathbb{R}^+$.

For notational convenience, we define the map

$$\Psi_t(A) := \cup_{(x,v)\in A} \psi_{(x,v)}(t) \tag{11}$$

for all $A \subset \mathcal{M}^\circ$ and $t > 0$.

We now prove that the stationary distribution of the geodesic walk is uniform with respect to the volume measure on any closed manifold $\mathcal{M}$; this proof will hold even if the curvature is not everywhere positive (Theorem 1). In Theorem 4 of Section 6 we prove that the stationary distribution is unique if $\mathcal{M}$ has bounded positive curvature.

Theorem 1. *The uniform distribution on $\mathcal{M}$ is a stationary distribution of the geodesic walk Markov chain defined in Algorithm 1.*

Proof. We begin by defining two transition kernels $\mathbb{K}_1^\circ$ and $\mathbb{K}_2^\circ$ on $\mathcal{M}^\circ$. We define the Markov chain $\{(X_i^{(1)}, v_i^{(1)})\}_{i\in\mathbb{N}}$ on $\mathcal{M}^\circ$ by the update rule

$$(X_{i+1}^{(1)}, v_{i+1}^{(1)}) = \psi_{(X_i^{(1)}, v_i^{(1)})}(T)$$

and define $\mathbb{K}_1^\circ$ to be the associated transition kernel. We then define the Markov chain $\{(X_i^{(2)}, v_i^{(2)})\}_{i\in\mathbb{N}}$ on $\mathcal{M}^\circ$ by the update rule

$$\begin{aligned} X_{i+1}^{(2)} &= X_i^{(2)} \\ v_{i+1}^{(2)} &\sim \mathrm{Unif}(\mathbb{S}(\mathcal{T}_{X_{i+1}^{(2)}})) \end{aligned}$$

and define $\mathbb{K}_2^\circ$ to be the associated transition kernel. We observe that $\mathbb{K}_1^\circ$ is deterministic, that $\mathbb{K}_2^\circ$ only ever updates its second coordinate, and finally that

$$\mathbb{K}^\circ = \mathbb{K}_2^\circ \mathbb{K}_1^\circ. \tag{12}$$

Since the Liouville measure is invariant under geodesic flow [11, 22, 3], we have

$$\mathcal{L}(A) = \mathcal{L}(\Psi_T(A)). \tag{13}$$

for every measurable $A \subset \mathcal{M}^\circ$. This implies that the Liouville measure $\mathcal{L}$ is invariant under $\mathbb{K}_1^\circ$:

$$\mathcal{L}\,\mathbb{K}_1^\circ = \mathcal{L}. \tag{14}$$

It is clear by inspection that

$$\mathcal{L}\,\mathbb{K}_2^\circ = \mathcal{L}. \tag{15}$$

Thus, by Equations 12, 14 and 15,

$$\mathcal{L}\,\mathbb{K}^\circ = \mathcal{L}. \tag{16}$$

Let $\{(X_i, U_i)\}_{i\geq 0}$ be a Markov chain started at $(X_1, U_1) \sim \mathcal{L}$ and evolving according to $\mathbb{K}^\circ$. By the definition of $\mathbb{K}^\circ$, the marginal process $\{X_i\}_{i\in\mathbb{N}}$ is a Markov chain evolving according to the transition kernel $\mathbb{K}$. Thus, the marginal distribution of $\mathcal{L}$ on its first coordinate $\mathcal{M}$ must

be a stationary distribution for $\mathbb{K}$. But this marginal distribution is exactly $\mathrm{Unif}(\mathcal{M})$, completing the proof. □

5. Coupling the Geodesic Walk

We define a coupling of two copies of the geodesic walk:

Definition 1 (Coupling of Geodesic Walk)**.** *Fix $x, y \in \mathcal{M}$ and $0 < T < \infty$. We define a pair of stochastic processes $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ with $X_1 = x$, $Y_1 = y$ as follows:*

Let $X_1 = x$ and $Y_1 = y$. For $i \in \mathbb{N}$, inductively sample

$$U_i \sim \mathrm{Unif}(\mathbb{S}(\mathcal{T}_{X_i}))$$

and set

$$V_i = \phi(U_i; X_i, Y_i) \tag{17}$$

$$X_{i+1} = \gamma_{(X_i,U_i)}(T) \tag{18}$$

$$Y_{i+1} = \gamma_{(Y_i,V_i)}(T). \tag{19}$$

This coupling is illustrated in Figure 2.

This stochastic process is a valid coupling of two copies of the geodesic walk:

Theorem 2. *Let $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ be as in Definition 1. Then the marginal processes $\{X_i\}_{i\in\mathbb{N}}$ and $\{Y_i\}_{i\in\mathbb{N}}$ are each Markov chains with transition kernel $\mathbb{K}$.*

Proof. It is straightforward to check that $\{X_i\}_{i\in\mathbb{N}}$ is exactly the Markov chain defined in Algorithm 1 (even the notation is the same). Thus, it remains only to check that $\{Y_i\}_{i\in\mathbb{N}}$ has the correct distribution.

To see that $\{Y_i\}_{i\in\mathbb{N}}$ has the correct distribution, looking at Algorithm 1 it is enough to check that (conditional on $\{Y_j\}_{j\le i}$), V_i has uniform distribution on $\mathbb{S}(\mathcal{T}_{Y_i})$.

By the fundamental theorem of Riemannian geometry, parallel transport using the Levi-Civita connection is an affine transformation. Hence, it preserves the angles between any two vectors u and u' parallel transported to $v = \phi(u; X_i, Y_i)$ and $v' = \phi(u'; X_i, Y_i)$, respectively:

$$\measuredangle(u, u') = \measuredangle(v, v'). \tag{20}$$

Therefore, denoting the conditional density of U_i conditioned on X_i by $f_{U_i}(u|X_i = x)$, and the conditional density of V_i conditioned on Y_i by $f_{V_i}(v|Y_i = y)$, we have:

$$f_{V_i}(v|Y_i = y) = f_{U_i}(u|X_i = x) = \frac{1}{\mathrm{Vol}(\mathbb{S}^{d-1})} \tag{21}$$

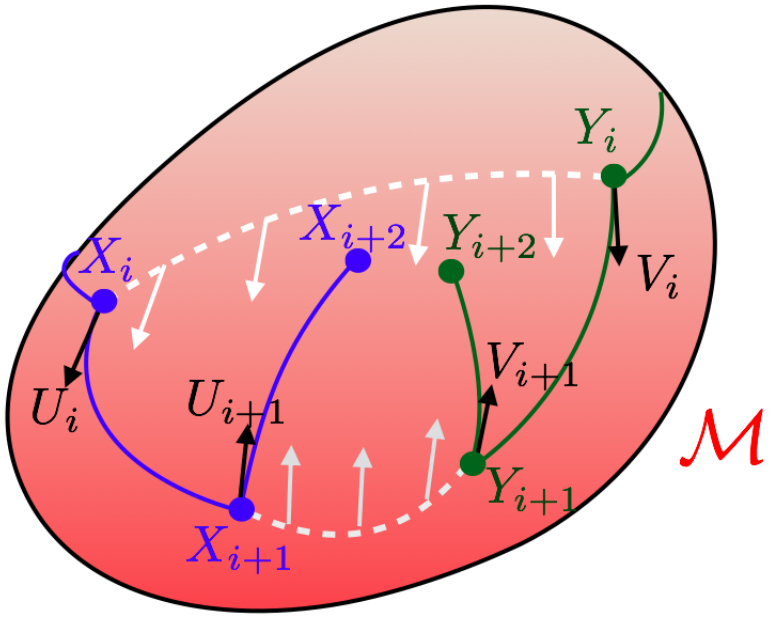

FIGURE 2. Coupling of the two geodesic walks $X_1, X_2, \dots$ (blue) and $Y_1, Y_2, \dots$ (green). The coupling is achieved by setting the initial trajectory velocity V_i at Y_i (black) to be the parallel transport of the initial trajectory velocity U_i at X_i along a minimum geodesic from X_i to Y_i (gray/white dashed line). Since $\mathcal{M}$ has positive curvature, the distance between the Markov chains contracts rapidly at each step.

whenever u is a vector on the unit sphere in the tangent plane of x, and $v = \phi(u; x, y)$ is its parallel transport to Y_i. Since the Levi-Civita connection also preserves magnitude, the fact that u is on the unit sphere in the tangent plane of x implies that v is on the unit sphere in the tangent plane of y.

Since U_i is uniformly distributed on the unit sphere in $\mathcal{T}_{X_i}$, this implies that V_i is uniformly distributed on the unit sphere in $\mathcal{T}_{Y_i}$. Therefore the transition kernel of the Markov chain $\{Y_i\}_{i\in\mathbb{N}}$ must be $\mathbb{K}$ as well. □

6. Contraction of Coupled Geodesics

In this section we prove a contraction bound on the coupled geodesic walks $X_1, X_2, \dots$ and $Y_1, Y_2, \dots$. The following extension of the Rauch comparison theorem from differential geometry will allow us to do so:

Lemma 1. *Let $x, y \in \mathcal{M}$, let $v_x \in \mathcal{T}_x$ and let $v_y = \phi(v_x; x, y)$. Then for all $0 \leq T \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$, we have*

$$\operatorname{dist}(\gamma_{(x,v_x)}(T), \gamma_{(y,v_y)}(T)) \leq \cos(\sqrt{\mathfrak{m}_2}T) \times \operatorname{dist}(x, y). \tag{22}$$

Remark 1. *An upper bound for the geodesic distance is proved in [9]. The authors mention that it is possible to prove a lower bound using a similar construction, although they do not do so explicitly. For this reason, we prove Lemma 1 explicitly here.*

Proof. Consider the family of geodesics $\{\gamma_{(\omega(\tau;x,y),\phi(\tau;v_x;x,y))}(t) : \tau \in [0, \operatorname{dist}(x, y)]\}$. Let $J_\tau(t)$ be the Jacobi field associated with this family.

Since $\mathcal{M}$ has curvature bounded above by $\mathfrak{M}_2 > 0$, the Rauch comparison theorem (comparing $\mathcal{M}$ to a sphere of radius $\frac{1}{\sqrt{\mathfrak{M}_2}}$) [17, 8] gives

$$\|J_\tau(t)\| \geq \cos(\sqrt{\mathfrak{M}_2}t) \tag{23}$$

for all

$$t \leq \inf\{t > 0 : \cos(\sqrt{\mathfrak{M}_2}t) = 0\} = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}.$$

Thus,

$$\|J_\tau(t)\| > 0 \tag{24}$$

for all $0 \leq t < T < \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$. Therefore, since $\|J_\tau(t)\|$ has no zeros on $[0, T)$, applying the Rauch comparison theorem a second time gives (this time comparing $\mathcal{M}$ to a sphere of radius $\frac{1}{\sqrt{\mathfrak{m}_2}}$):

$$\|J_\tau(t)\| \leq \cos(\sqrt{\mathfrak{m}_2}t) \tag{25}$$

for all $0 \leq t < T$.

Define

$$F(\tau) := \operatorname{length}(\{\gamma_{(\omega(s;x,y),\phi(s;v_x;x,y))}(T) : s \in [0, \tau]\}) \tag{26}$$

for $0 < \tau < \operatorname{dist}(x, y)$. Then

$$\frac{\mathrm{d}F}{\mathrm{d}\tau} = \|J_\tau(T)\|. \tag{27}$$

Therefore, by the fundamental theorem of calculus,

$$\begin{aligned} F(\operatorname{dist}(x, y)) &= \int_0^{\operatorname{dist}(x,y)} \frac{\mathrm{d}F}{\mathrm{d}\tau} \mathrm{d}\tau \\ &\overset{\text{Eq. 27}}{=} \int_0^{\operatorname{dist}(x,y)} \|J_\tau(T)\| \mathrm{d}\tau \\ &\overset{\text{Eq. 25}}{\leq} \int_0^{\operatorname{dist}(x,y)} \cos(\sqrt{\mathfrak{m}_2}T) d\tau \end{aligned}$$

$$= \operatorname{dist}(x,y) \times \cos(\sqrt{\mathfrak{m}_2}T). \tag{28}$$

□

Theorem 3. *Fix $T \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$ and let $\{(X_i, Y_i)\}_{i\geq 0}$ be coupled as in Definition 1. Then*

$$\operatorname{dist}(X_{i+1}, Y_{i+1}) \leq \operatorname{dist}(X_i, Y_i) \times \cos(\sqrt{\mathfrak{m}_2}T). \tag{29}$$

Proof. This follows immediately from an application of Lemma 1. □

Theorem 4. *The uniform distribution on $\mathcal{M}$ is the unique stationary distribution of the geodesic walk Markov chain defined in Algorithm 1.*

Proof. By Theorem 1, the uniform distribution Π on $\mathcal{M}$ is a stationary distribution of $\mathbb{K}$. Assume that there is a second stationary distribution $\Pi' \neq \Pi$. Since $\mathcal{M}$ has finite diameter, this implies $0 < W_d(\Pi, \Pi') < \infty$. By Theorem 3, however,

$$W_d(\Pi, \Pi') = W_d(\Pi\,\mathbb{K}, \Pi'\,\mathbb{K}) \leq \cos(\sqrt{\mathfrak{m}_2}T)\, W_d(\Pi, \Pi') < W_d(\Pi, \Pi').$$

Since $0 < W_d(\Pi, \Pi') < \infty$, this is a contradiction, and so no such distribution Π' exists. □

7. Bounding the Mixing Time

In this section we bound the mixing time (in Wasserstein Transportation distance) using a type of argument similar to those used in [14].

Theorem 5. *The Wasserstein mixing profile $\mathfrak{t}_{\text{mix}}$ of the Markov chain with transition kernel $\mathbb{K}$ defined in Section 2 satisfies*

$$\mathfrak{t}_{\text{mix}}(\epsilon) \leq \left\lceil \frac{\log(\epsilon D^{-1})}{\log(\cos(\sqrt{\mathfrak{m}_2}T))} \right\rceil \right\rceil = \mathcal{O}\left(\frac{1}{\mathfrak{m}_2 T^2} \cdot \log(D\epsilon^{-1}) \right). \tag{30}$$

In particular, for $T = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$, the mixing time is bounded by

$$\mathfrak{t}_{\text{mix}}(\epsilon) \leq \left\lceil \frac{\log(\epsilon D^{-1})}{\log(\cos(\frac{\pi}{2}\frac{\sqrt{\mathfrak{m}_2}}{\sqrt{\mathfrak{M}_2}}))} \right\rceil = \mathcal{O}\left(\frac{\mathfrak{M}_2}{\mathfrak{m}_2} \cdot \log(D\epsilon^{-1}) \right). \tag{31}$$

Proof. Fix $0 \leq T \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$. Let $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ be two copies of the Markov chain with kernel $\mathbb{K}$, coupled as in Definition 1. Recall that, by Theorem 2, this is a valid coupling. By Theorem 3, this chain satisfies the following contraction inequality:

$$\operatorname{dist}(X_{i+1}, Y_{i+1}) \leq \operatorname{dist}(X_i, Y_i) \times \cos(\sqrt{\mathfrak{m}_2}T) \qquad \forall i > 0. \tag{32}$$

Applying this contraction inequality repeatedly, we find:

(33)
$$\mathrm{dist}(X_j, Y_j) \leq \mathrm{dist}(X_0, Y_0) \times \cos(\sqrt{\mathfrak{m}_2} T)^j \leq D \times \cos(\sqrt{\mathfrak{m}_2} T)^j \qquad \forall j > 0.$$

This bound immediately implies that, for all $0 < \epsilon < 1$,

$$\mathfrak{t}_{\mathrm{mix}}(\epsilon) \leq \left\lceil \frac{\log(\epsilon D^{-1})}{\log(\cos(\sqrt{\mathfrak{m}_2} T))} \right\rceil. \tag{34}$$

In particular, for $T = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$, we get:

$$\mathfrak{t}_{\mathrm{mix}}(\epsilon) \leq \left\lceil \frac{\log(\epsilon D^{-1})}{\log(\cos(\pi \frac{\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \right\rceil. \tag{35}$$

□

8. Approximating Geodesics

In this section we describe and analyze an approximation to the geodesic walk, showing that it can be used to approximately sample from the uniform distribution on $\mathcal{M}$ to arbitrary accuracy using a dimension-independent number of computations.

We begin in Section 8.1 by assuming that we have access to an oracle ★, described at the beginning of Section 8.1, that can approximate short geodesic paths with small error. We use this oracle to construct an "approximate" version of the geodesic walk. Finally, we prove that, for any fixed error rate $\epsilon > 0$, this approximate geodesic walk can give samples that are within ϵ of the uniform distribution on $\mathcal{M}$ in Wasserstein distance using a dimension-independent number of oracle calls. In Section 8.2, we construct ★ explicitly for the special case where $\mathcal{M}$ is the boundary of a convex body, using only the basic convex body oracles used in [7].

8.1. Approximating Geodesics on General Positive-Curvature Manifolds.

In Algorithm 2 we compute recursively an approximation $\gamma^{\dagger}_{\theta;(x,v)}(T)$ of $\gamma_{(x,v)}(T)$ by repeatedly calling an oracle ★. In particular, we will assume the following about the oracle ★:

Definition 2. *The oracle* $\bigstar : \mathcal{M}^\circ \times \mathbb{R}^+ \to \mathcal{M}^\circ \times \mathbb{R}^+$ *is a function that has inputs* $(x, v) \in \mathcal{M}^\circ$*, and* $\theta > 0$*, and outputs* $\psi^{\star}_{\theta;(x,v)} = \left(\gamma^{\star}_{\theta;(x,v)}, \varphi^{\star}_{\theta;(x,v)}\right) \in \mathcal{M}^\circ$ *and* $\Delta^{\star}_{\theta;(x,v)} \in \mathbb{R}^+$.

The oracle tries to approximate the geodesic trajectory $\psi_{(x,v)}(t)$ at $t = \Delta^{\star}_{\theta;(x,v)}$, where the step size $\Delta^{\star}_{\theta;(x,v)}$ is determined by the oracle.

The parameter $\theta > 0$ allows the user to adjust the accuracy of the oracle's approximation.

We also define for every $\theta > 0$,

$$\Delta^\star_{\theta;\min} := \inf_{(x,v)\in\mathcal{M}^\circ} \Delta^\star_{\theta;(x,v)}. \tag{36}$$

In particular, we will assume the following about the oracle ★:

Assumption 1. *For every $\theta > 0$, there exist multivariate polynomials P and Q with the property that, at every $(x, v) \in \mathcal{M}^\circ$, we have*

$$\frac{1}{\sqrt{\mathfrak{M}_2}} \cdot P\left(\theta, \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}, \frac{\sqrt{\mathfrak{m}_2}}{\sqrt{\mathfrak{M}_2}}\right) \leq \Delta^\star_{\theta;(x,v)} \leq \frac{1}{\sqrt{\mathfrak{M}_2}} \cdot Q\left(\theta, \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}, \frac{\sqrt{\mathfrak{m}_2}}{\sqrt{\mathfrak{M}_2}}\right). \tag{37}$$

Also, $P(\theta, \cdot, \cdot)$ and $Q(\theta, \cdot, \cdot)$ are strictly increasing in θ for all $\theta > 0$ when the other two variables are positive, and $P(0, \cdot, \cdot) = Q(0, \cdot, \cdot) = 0$.

Assumption 2. *There exist constants $\alpha > 0$, $\beta > 0$ such that*

$$\operatorname{dist}(\gamma_{(x,v)}(\Delta^\star_{\theta;(x,v)}), \gamma^\star_{\theta;(x,v)}) \leq \alpha \cdot \theta \cdot \Delta^\star_{\theta;(x,v)} \tag{38}$$

and also

$$\|\varphi^\star_{\theta;(x,v)} - \overline{\varphi_{(x,v)}(\Delta^\star_{\theta;(x,v)})}\| \leq \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta^\star_{\theta;(x,v)} \tag{39}$$

for every $(x, v) \in \mathcal{M}^\circ$, where $\overline{\varphi_{(x,v)}(\Delta^\star_{\theta;(x,v)})} := \phi(\varphi_{(x,v)}(\Delta^\star_{\theta;(x,v)}); \gamma_{(x,v)}(\Delta^\star_{\theta;(x,v)}), \gamma^\star_{(x,v)})$.

For the rest of Section 8.1, we will assume that the oracle ★ satisfies Assumptions 1 and 2.

Algorithm 2 Numerical Approximation of Geodesic Trajectory

parameters: $\theta > 0$, $T > 0$.
input: Manifold $(\mathcal{M}, g)$ with sectional curvature bounded above by $\mathfrak{M}_2$.
input: Initial point $x_0 \in \mathcal{M}$, initial velocity $v_0 \in \mathcal{T}_{x_0}\mathcal{M}$.
input: Oracle $\bigstar$ returning $\psi^{\star}_{\theta;(x,v)}$ and $\Delta^{\star}_{\theta;(x,v)}$ for any $(x,v) \in \mathcal{M}^{\circ}$.
output: $\gamma^{\dagger}_{\theta;(x,v)}(T)$.

1: Set $i = 0$, $\Delta_0 = 0$.
2: **while** $\sum_{j=0}^{i} \Delta_j \leq T$ **do**
3: Call the oracle $\bigstar$ to compute $(x^{\dagger}_{i+1}, v^{\dagger}_{i+1}) = \psi^{\star}_{\theta;(x_i,v_i)}$ and $\Delta_{i+1} = \Delta^{\star}_{\theta;(x^{\dagger}_i, v^{\dagger}_i)}$.
4: Set $i = i + 1$.
5: **end while**
6: Set $\delta = T - \sum_{j=1}^{i} \Delta_j$.
7: Using the bisection method, iteratively call $\bigstar$ to find a value $\hat{\theta}$ for which $\delta - \frac{\theta}{\sqrt{\mathfrak{M}_2}} \leq \Delta^{\star}_{\hat{\theta};(x^{\dagger}_i, v^{\dagger}_i)} \leq \delta$. Set $\Delta_{i+1} := \Delta^{\star}_{\hat{\theta};(x^{\dagger}_i, v^{\dagger}_i)}$.
8: Call the oracle $\bigstar$ to compute $x^{\dagger}_{i+1} = \gamma^{\star}_{\hat{\theta};(x_i,v_i)}$.
9: Set $\gamma^{\dagger}_{\theta;(x,v)}(T) = x^{\dagger}_{i+1}$.

Using the notation of Algorithm 1, the geodesic walk Markov chain is defined recursively by

$$X_{i+1} := \gamma_{(X_i,U_i)}(T) \qquad \forall i = 0, 1, 2, \ldots \tag{40}$$

where U_i is uniformly distributed on $\mathbb{S}(\mathcal{T}_{X_i})$. For fixed $\theta > 0$, we can use Algorithm 2 to recursively generate an approximate geodesic walk $X^{\theta}_1, X^{\theta}_2, \ldots$ coupled to the geodesic walk $X_1, X_2, \ldots$, using the following recursion:

$$X^{\theta}_0 = X_0 \tag{41}$$

$$X^{\theta}_{i+1} = \gamma^{\dagger}_{\theta;(X^{\theta}_i, U'_i)}(T) \qquad \forall i = 0, 1, 2, \ldots, \tag{42}$$

where $U'_i = \phi(U_i; X_i, X^{\theta}_i)$. For $i \in \mathbb{N}$, we let Π^{θ}_i be the distribution of X^{θ}_i.

Lemma 2. *For every $(x_0, v_0) \in \mathcal{M}^{\circ}$ and every $T \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$, Algorithm 2 makes at most $\lceil \frac{T}{\Delta_{\theta;\min}} \rceil + \mathcal{O}(\log(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}))$ calls of the oracle $\bigstar$ and returns a point $\gamma^{\dagger}_{\theta;(x_0,v_0)}(T)$ such that*

$$\operatorname{dist}(\gamma_{(x_0,v_0)}(T), \gamma^{\dagger}_{\theta;(x_0,v_0)}(T)) \leq \left[1 + \frac{\pi}{2}\alpha + (\frac{\pi}{2})^2\beta\right] \cdot \frac{1}{\sqrt{\mathfrak{M}_2}} \cdot \theta \tag{43}$$

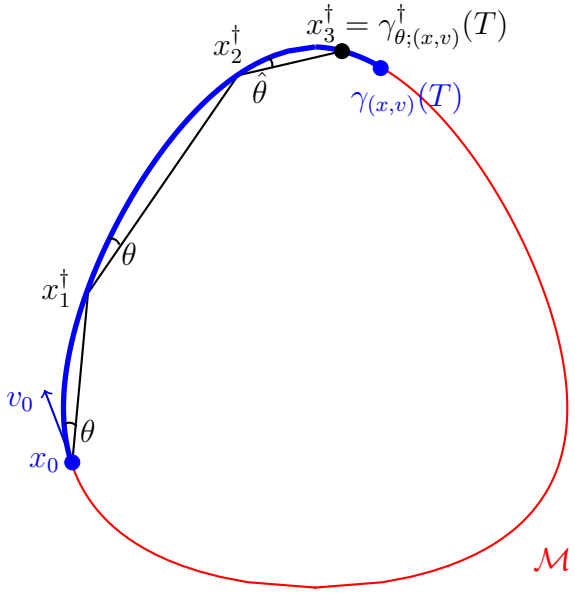

FIGURE 3. Algorithm for approximating geodesics. In this example $\mathcal{M}$ is the boundary of a convex body, so one can generate each step in the approximation as a line ℓ_i leaving the current approximation point $x_i^\dagger$ on the boundary at a fixed angle θ. The next approximation point $x_{i+1}^\dagger$ is the next point when ℓ_i intersects the boundary. Lemmas 2 and 6 show that the final point $x_{i_{\max}+1}^\dagger = \gamma^\dagger_{\theta;(x,v)}(T)$ (here $i_{\max} = 2$) approximates the final point $\gamma_{(x,v)}(T)$ of the true geodesic trajectory to arbitrary accuracy ϵ in $i_{\max}(\epsilon)$ steps, where $i_{\max}(\epsilon)$ is independent of the dimension, if θ is appropriately chosen.

Proof. (1) **Counting the number of ★ Oracle calls:** Define $i_{\max} := \max\{i \in \mathbb{N} : \sum_{j=0}^{i+1} \Delta_j \leq T\}$. The number of ★ oracle calls required for Step 3 in Algorithm 2 is at most $\lceil \frac{T}{\Delta_{\theta;\min}} \rceil$, since the largest possible value of i such that $\sum_{j=1}^{i+1} \Delta_j \leq T$ is $\lfloor \frac{T}{\Delta_{\theta;\min}} \rfloor$. Assumption 1 implies that the number of ★ oracle calls the bisection method makes in Step 7 of Algorithm 2 to find a value $\hat{\theta}$ for which $\Delta_{i+1} := \Delta^\star_{\hat{\theta};(x_i^\dagger, v_i^\dagger)}$ satisfies $\delta - \frac{\theta}{\sqrt{\mathfrak{M}_2}} \leq \Delta_{i+1} \leq \delta$ is at most $\mathcal{O}(\log(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}))$. Step 8 makes exactly one call of the oracle ★.

(2) **Bounding the error:** Let $x_i^\dagger$, $v_i^\dagger$, Δ_i, θ, and $\hat{\theta}$ be as in Algorithm 2.

Define $(x_i, v_i) := \psi_{(x_{i-1}^\dagger, v_{i-1}^\dagger)}(\Delta_i)$, and $\bar{v_i} := \phi(v_i; x_i, x_i^\dagger)$. By Inequality 39, we have $\|v_i^\dagger - \bar{v_i}\| \leq \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta_i$.

Therefore, by the Toponogov triangle comparison theorem (comparing to a triangle in a sphere of radius $\frac{1}{\mathfrak{M}_2}$), we have

$$\operatorname{dist}(\gamma_{(x_i^\dagger, \bar{v_i})}(t), \gamma_{(x_i^\dagger, v_i^\dagger)}(t)) \leq t \cdot \|v_i^\dagger - \bar{v_i}\| \leq t \cdot \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta_i \tag{44}$$

for all $t > 0$.

By Inequality 38 we have $\operatorname{dist}(x_i, x_i^\dagger) \leq \alpha \cdot \theta \cdot \Delta_i$. Therefore, by Lemma 1,

$$\operatorname{dist}(\gamma_{(x_i, v_i)}(t), \gamma_{(x_i^\dagger, \bar{v_i})}(t)) \leq \operatorname{dist}(x_i, x_i^\dagger) \leq \alpha \cdot \theta \cdot \Delta_i \tag{45}$$

for all $0 \leq t \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$.

Define the "remaining time after the i'th step" to be $T_i := \sum_{j=i+1}^{i_{\max}+1} \Delta_j$. Then

$$\begin{aligned}
\operatorname{dist}(\gamma_{(x_0,v_0)}(T), \gamma^\dagger_{(x_0,v_0)}(T)) &\leq \operatorname{dist}(\gamma_{(x_0,v_0)}(T), \gamma_{(x_0,v_0)}(T_0)) + \operatorname{dist}(\gamma_{(x_0,v_0)}(T_0), x^\dagger_{i_{\max}}) \\
&\leq T - T_0 + \operatorname{dist}(\gamma_{(x_0,v_0)}(T_0), x^\dagger_{i_{\max}}) \\
&\leq T - T_0 + \sum_{i=1}^{i_{\max}+1} (\operatorname{dist}(\gamma_{(x_i,v_i)}(T_i), \gamma_{(x_i^\dagger,\bar{v_i})}(T_i)) + \operatorname{dist}(\gamma_{(x_i^\dagger,\bar{v_i})}(T_i), \gamma_{(x_i^\dagger,v_i^\dagger)}(T_i))) \\
&\leq \frac{\theta}{\sqrt{\mathfrak{M}_2}} + \sum_{i=1}^{i_{\max}+1} (\alpha \cdot \theta \cdot \Delta_i + T_i \cdot \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta_i) \\
&\leq \frac{\theta}{\sqrt{\mathfrak{M}_2}} + [\alpha \cdot \theta + T_0 \cdot \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta] \cdot \sum_{i=0}^{i_{\max}+1} \Delta_i \\
&\leq \frac{\theta}{\sqrt{\mathfrak{M}_2}} + [\alpha \cdot \theta + T_0 \cdot \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta] \cdot T_0 \\
&\leq \frac{\theta}{\sqrt{\mathfrak{M}_2}} + [\alpha + \frac{\pi}{2\sqrt{\mathfrak{M}_2}} \cdot \beta \cdot \sqrt{\mathfrak{M}_2}] \cdot \frac{\pi}{2\sqrt{\mathfrak{M}_2}} \cdot \theta, \\
&= \left[1 + \frac{\pi}{2}\alpha + (\frac{\pi}{2})^2 \beta\right] \cdot \frac{1}{\sqrt{\mathfrak{M}_2}} \cdot \theta.
\end{aligned}$$

The second inequality holds because $\gamma_{(x_0,v_0)}(t)$ has unit velocity for every $t \in \mathbb{R}$ and in particular for every $t \in [T_0, T]$. The third inequality is the triangle inequality (see Figure 4), the fourth inequality is an application of Equations 44 and 45, and the last inequality uses our assumption that $T_0 \leq T \leq \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$.

□

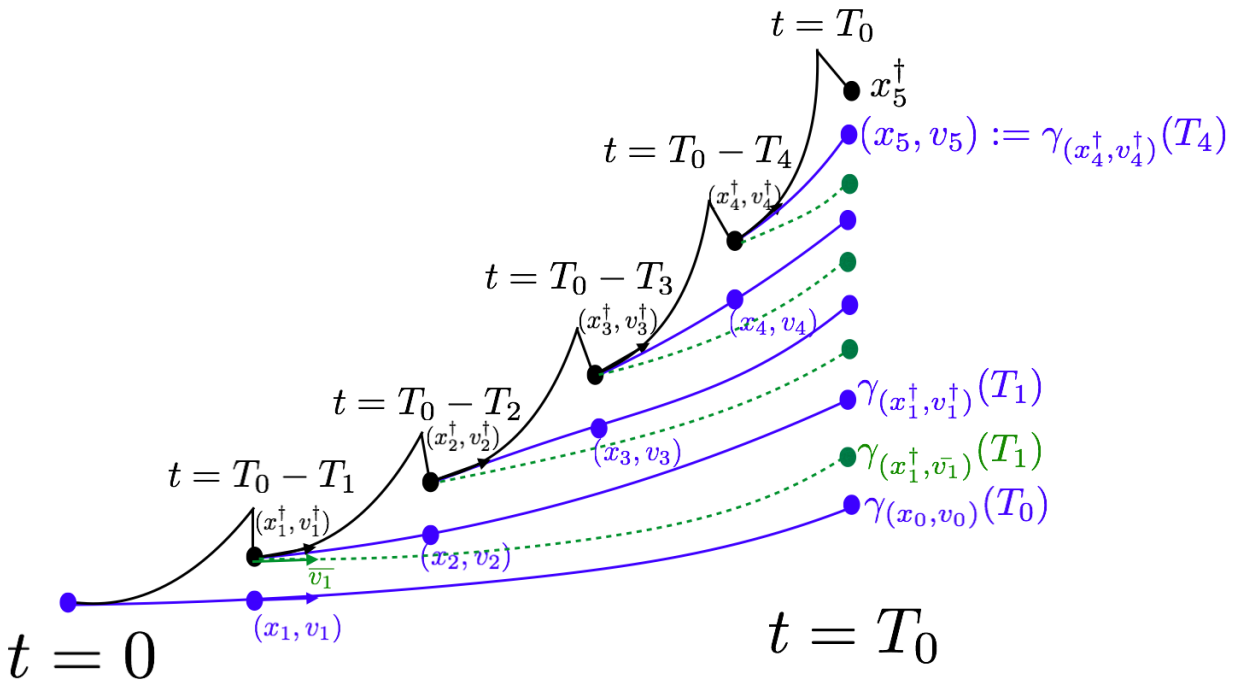

FIGURE 4. This is an illustration of the proof of Lemma 2. Steps taken by the oracle ★ (which is illustrated here as going along a short curved path and then projecting its position $x_i^\dagger := \gamma^\star_{(x_{i-1}^\dagger, v_{i-1}^\dagger)}$ (black dot) and velocity $v_i^\dagger := \gamma^\star_{(x_{i-1}^\dagger, v_{i-1}^\dagger)}$ (black arrow) back onto the manifold) are in black. The true geodesic paths are blue curves or green dashed curves. Only the geodesic path $\gamma_{(x_0,v_0)}(t)$ on the bottom belongs to the true geodesic walk. We imagine the other geodesic paths to help us bound the error. The distance between the blue dot x_i and black dot $x_i^\dagger$ at each time $t = T_0 - T_i$, and the angle between the velocities v_i and $\bar{v}_i$, where v_i is the velocity at the blue dot and $\bar{v}_i := \phi(v_i, x_i, x_i^\dagger)$, are bounded because of our assumptions on the accuracy of the oracle ★. The distance from any green dot at $t = T_0$ to the blue dot directly below it is bounded using Lemma 1. The distance from that same green dot to the blue dot directly above it is bounded using the Toponogov triangle comparison theorem.

We show that the random walk described by Algorithm 2 can be used to sample approximately uniformly from $\mathcal{M}$. We first need a simple probabilistic bound:

Lemma 3. *Let K be a transition kernel on metric space (Ω, d) with unique stationary measure μ and contraction coefficient $\kappa > 0$ satisfying*

$$W_d(K(x,\cdot), K(y,\cdot)) \leq (1-\kappa)d(x,y) \tag{46}$$

for all $x, y \in \Omega$. Let Q be a transition kernel on (Ω, d) with stationary measure ν. Assume that

$$\sup_{x\in\Omega} W_d(K(x,\cdot), Q(x,\cdot)) < \delta \tag{47}$$

for some fixed $\delta \geq 0$. Also assume that (Ω, d) has diameter $0 < D < \infty$. Then Q satisfies

$$W_d(Q^t(x,\cdot), \mu) \leq (1-\kappa)^t D + \frac{\delta}{\kappa} \tag{48}$$

for all $x \in \Omega$ and $t \in \mathbb{N}$, and furthermore

$$W_d(\mu, \nu) \leq \frac{\delta}{\kappa}. \tag{49}$$

Proof. Fix a tolerance $\eta > 0$ and starting points $x, y \in \Omega$. We begin by constructing a coupling of $Q(x,\cdot)$ and $K(y,\cdot)$.

By Inequality 46, it is possible to couple two random variables $X' \sim K(x,\cdot)$ and $Y' \sim K(y,\cdot)$ so that

$$\mathbb{E}[d(X', Y')] \leq (1-\kappa)d(x,y) + \eta. \tag{50}$$

Furthermore, by Inequality 47 and the standard gluing lemma (see Chapter 1 of [21]), it is possible to couple X', Y' to the random variable $X \sim Q(x,\cdot)$ so that

$$\mathbb{E}[d(X, X')] \leq \delta + \eta. \tag{51}$$

Combining Inequalities 50 and 51, it is possible to couple $X \sim Q(x,\cdot)$ and $Y' \sim K(y,\cdot)$ so that

$$\mathbb{E}[d(X, Y')] \leq \mathbb{E}[d(X, X')] + \mathbb{E}[d(X', Y')] \tag{52}$$

$$\leq \delta + (1-\kappa)d(x,y) + 2\eta. \tag{53}$$

We denote by M the kernel on Ω^2 given by this coupling of $Q(x,\cdot)$ and $K(y,\cdot)$.

Let $x \in \Omega$, and let $Y \sim \mu$. Let $\{(X_t, Y_t)\}_{t \geq 0}$ be a Markov chain evolving according to the kernel M with initial conditions $(X_0, Y_0) = (x, Y)$. By Inequality 52, we have for all $t \in \mathbb{N}$ that

$$\mathbb{E}[d(X_t, Y_t)] \leq \mathbb{E}[(1-\kappa)\mathbb{E}[d(X_{t-1}, Y_{t-1})] + \delta + 2\eta] \tag{54}$$

$$\leq (1-\kappa)^2 \mathbb{E}[d(X_{t-2}, Y_{t-2})] + (\delta + 2\eta)(1 + (1-\kappa)) \tag{55}$$

$$\leq \dots \tag{56}$$

$$\leq (1-\kappa)^t \mathbb{E}[d(X_0, Y_0)] + \frac{\delta + 2\eta}{\kappa}. \tag{57}$$

Since this holds for arbitrary $\eta > 0$ and $Y_t \sim \mu$ for all $t \in \mathbb{N}$, we have

$$W_d(Q^t(x, \cdot), \mu) \leq (1-\kappa)^t D + \frac{\delta}{\kappa}. \tag{58}$$

This completes the proof of Inequality 48. Inequality 49 follows immediately from letting t go to infinity. □

We apply this bound:

Lemma 4. *Let $\epsilon > 0$. Set $T = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$ and set*

$$\theta(\epsilon) = \epsilon \cdot \sqrt{\mathfrak{M}_2} \cdot \frac{1 - \cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}})}{2(1 + \frac{\pi}{2}\alpha + (\frac{\pi}{2})^2\beta)} \tag{59}$$

$$I(\epsilon) = \frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))}. \tag{60}$$

Then the distribution $\Pi_i^{\theta(\epsilon)}$ of the i'th step of the chain described in Equation 41 satisfies

$$\mathrm{Wass}(\Pi, \Pi_i^{\theta(\epsilon)}) \leq \epsilon \tag{61}$$

for all $i > I(\epsilon)$.

Proof. We apply Lemma 3, with $K = \mathbb{K}$ and Q the kernel of the Markov chain in Equation 41 with parameter $\theta = \theta(\epsilon)$ and $T = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$. We keep the notation of Lemma 3 and check that the assumptions are satisfied.

By Theorem 3, Inequality 46 is satisfied with

$$\kappa = 1 - \cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}).$$

By Lemma 2, Inequality 47 is satisfied with

$$\delta = \frac{1 + \frac{\pi}{2}\alpha + (\frac{\pi}{2})^2\beta}{\sqrt{\mathfrak{M}_2}} \cdot \theta.$$

Finally, by Inequality 8, the diameter of $\mathcal{M}$ satisfies

$$D \leq \frac{\pi}{\sqrt{\mathfrak{m}_2}}.$$

Applying Lemma 3 with these values of κ, δ and D completes the proof of this theorem.

□

This immediately gives the following dimension-free bound on the number of oracle calls required to obtain a sample from the uniform distribution on $\mathcal{M}$ with error less than some fixed $\epsilon > 0$:

Theorem 6. *Fix $\epsilon > 0$ and let $\theta(\epsilon), I(\epsilon)$ be as in Equation 59. Also let $X_{I(\epsilon)+1}^{\theta(\epsilon)}$ be the $(I(\epsilon)+1)$'st point generated by Equation 41 with parameters $\theta = \theta(\epsilon)$, $T = \frac{\pi}{2\sqrt{\mathfrak{M}_2}}$.*

Then the distribution $\Pi_{I(\epsilon)+1}^{\theta(\epsilon)}$ of $X_{I(\epsilon)+1}^{\theta(\epsilon)}$ satisfies

$$W_d(\Pi_{I(\epsilon)+1}^{\theta(\epsilon)}, \Pi) \leq \epsilon, \tag{62}$$

and the number $N(\epsilon)$ of ★ oracle calls used to generate $X_{I(\epsilon)+1}^{\theta(\epsilon)}$ satisfies

$$N(\epsilon) \leq \frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \left(\lceil \frac{\pi}{2\sqrt{\mathfrak{M}_2}\,\Delta_{\theta(\epsilon);\min}} \rceil + \mathcal{O}(\log(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}})) \right)$$

$$= \mathcal{O}^*\left(\frac{\mathfrak{M}_2}{\mathfrak{m}_2} \times \frac{1}{\sqrt{\mathfrak{M}_2}\Delta_{\theta(\epsilon);\min}} \right). \tag{63}$$

Proof. Inequality 62 follows immediately from Lemma 4.

By Lemma 2,

$$N(\epsilon) \leq I(\epsilon) \times \left(\lceil \frac{T}{\Delta_{\theta(\epsilon);\min}} \rceil + \mathcal{O}(\log(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}})) \right)$$

$$\leq \frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \left(\lceil \frac{\pi}{2\sqrt{\mathfrak{M}_2}\,\Delta_{\theta(\epsilon);\min}} \rceil + \mathcal{O}(\log(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}})) \right).$$

This completes the proof of the Theorem.

□

8.2. **Approximating Geodesics on Convex body Boundaries.** In this section we show how to construct an oracle that satisfies the requirements of Lemma 2 in the important special case when $\mathcal{M} = \partial\mathcal{K}$ is the boundary of a convex body $\mathcal{K}$. To build our oracle ★ (Algorithm 3), we use only the same basic inclusion-exclusion oracle and oracle for

the normal line to $\partial\mathcal{K}$ used in the stochastic billiards algorithm of [7]. In particular, one can construct an intersection oracle that returns the intersection of any line with $\mathcal{K}$ using the bisection method and the inclusion-exclusion oracle, which we will use to build $\bigstar$.

In the remainder of the paper, for any line segment $\ell_{a,b} \subset \mathbb{R}^{d+1}$ connecting points $a, b \in \mathbb{R}^{d+1}$, we denote its Euclidean length by $\text{length}(\ell_{a,b}) := \|b - a\|$.

Algorithm 3 An oracle $\bigstar$ in the special case where $\mathcal{M}$ is boundary of convex body $\mathcal{K}$

input: Intersection oracle for the convex body $\mathcal{K}$.
input: Oracle for the normal vector of $\mathcal{M} = \partial\mathcal{K}$.
input: $x \in \mathcal{K}$, $v \in \mathcal{T}_x\mathcal{K}$, $\theta > 0$.
output: $\Delta^\star_{\theta;(x,v)}$, $\psi^\star_{\theta;(x,v)} := (\gamma^\star_{\theta;(x,v)}, \varphi^\star_{\theta;(x,v)})$.

1: Generate the unique line ℓ through x at an angle θ to the tangent plane $\mathcal{T}_x$ whose projection onto $\mathcal{T}_x$ is parallel to v (using the normal vector oracle).
2: solve for the other intersection point x_1 in $\ell \cap \mathcal{M}$ by calling the intersection oracle.
3: set v_1 to be the projection of ℓ onto the tangent space $\mathcal{T}_{x_1}$ at x_1.
4: set $(\gamma^\star_{\theta;(x,v)}, \varphi^\star_{\theta;(x,v)}) = (x_1, v_1)$.
5: set $\Delta^\star_{\theta;(x,v)} = \|x_1 - x\|$.

In this section we will make a slightly strengthened assumption about the curvature. Towards this end, we define the inner radius of curvature at a point $x \in \partial\mathcal{K}$ to be the radius of the largest sphere in $\mathbb{R}^{d+1}$ that is tangent to $\partial\mathcal{K}$ at x and contained in $\mathcal{K}$. Similarly, we define the outer radius of curvature to be radius of the smallest sphere that is tangent to $\partial\mathcal{K}$ at x that contains $\mathcal{K}$.

Assumption 3. *$\mathcal{M} = \partial\mathcal{K}$ has inner radius of curvature uniformly bounded below by $\frac{1}{\sqrt{\mathfrak{M}_2}}$ and outer radius of curvature uniformly bounded above by $\frac{1}{\sqrt{\mathfrak{m}_2}}$.*

In particular, the above assumptions on the inner and outer radii of curvature imply our previous assumptions that the scalar curvature is bounded above and below by $\mathfrak{M}_2$ and $\mathfrak{m}_2$, respectively, although the converse is not true.

The following Lemma (Lemma 5) relates the geodesic distance on a convex body to the Euclidean distance in the ambient space. We will use Lemma 5 to prove the main result of this section (Lemma 6).

Before we state Lemma 5, we define $\text{dist}_{\mathcal{N}}(x, y)$ to be the geodesic distance between two points $x, y \in \mathcal{N}$ in a manifold $\mathcal{N}$. For every set $S = \{x, y\}$ consisting of two distinct points $x, y \in \mathcal{N}$, we also define $\text{dist}_{\mathcal{N}}(S) := \text{dist}_{\mathcal{N}}(x, y)$.

Lemma 5. *Let $q^{(1)}$ and $q^{(2)}$ be two points on the boundary $\partial\mathcal{K}$ of a convex body $\mathcal{K}$, with inner and outer radii of curvature bounded below and above by $\frac{1}{\sqrt{\mathfrak{M}_2}}$ and $\frac{1}{\sqrt{\mathfrak{m}_2}}$, respectively. Then*

$$\text{dist}_{\partial\mathcal{K}}(q^{(1)}, q^{(2)}) \leq 2\pi \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} \|q^{(2)} - q^{(1)}\| \tag{64}$$

Proof. Let ℓ be the line passing through both $q^{(1)}$ and $q^{(2)}$. Let B and B' be balls tangent to $\partial\mathcal{K}$ at $q^{(2)}$ of radius $\frac{1}{\sqrt{\mathfrak{M}_2}}$ and $\frac{1}{\sqrt{\mathfrak{m}_2}}$, respectively, with

$$B \subseteq \mathcal{K} \subseteq B'. \tag{65}$$

Since B, $\mathcal{K}$, and B' are all convex, we therefore have that (Figure 5),

$$\text{dist}_{\partial\mathcal{K}}(\ell \cap \partial\mathcal{K}) \leq \text{dist}_{\partial\mathcal{K}\cap A}(\ell \cap \partial\mathcal{K}) \leq \text{dist}_{\partial B'}(\ell \cap \partial B') + \text{length}(\ell \cap B') \leq 2\text{dist}_{\partial B'}(\ell \cap \partial B') \tag{66}$$

where A is the 2-plane containing the great circle (i.e. the spherical geodesic) in B' that connects the two points in $\ell \cap \partial B'$.

Since the ratio $\frac{\text{dist}_{\partial B'}(\ell \cap \partial B')}{\text{length}(\ell \cap B')}$ is maximized when $\ell \cap B'$ is the diameter of B', it must be true that

$$\text{dist}_{\partial B'}(\ell \cap B') \leq \pi \text{length}(\ell \cap B'). \tag{67}$$

Defining $q^{(2)}$ to be the origin, we get that

$$B' = \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} B. \tag{68}$$

So,

$$\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}(\ell \cap B) = \ell \cap (\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} B) = \ell \cap B'. \tag{69}$$

Hence,

$$\text{length}(\ell \cap B') = \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} \text{length}(\ell \cap B) = \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} \|q^{(2)} - q^{(1)}\|. \tag{70}$$

Therefore, combining Equations 66, 67, and 70, we get:

$$\text{dist}_{\partial\mathcal{K}}(q^{(1)}, q^{(2)}) = \text{dist}_{\partial\mathcal{K}}(\ell \cap \partial\mathcal{K}) \leq 2\text{dist}_{\partial B'}(\ell \cap \partial B')$$

$$\leq \pi \text{length}(\ell \cap B') = 2\pi \frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}} \|q^{(2)} - q^{(1)}\|. \tag{71}$$

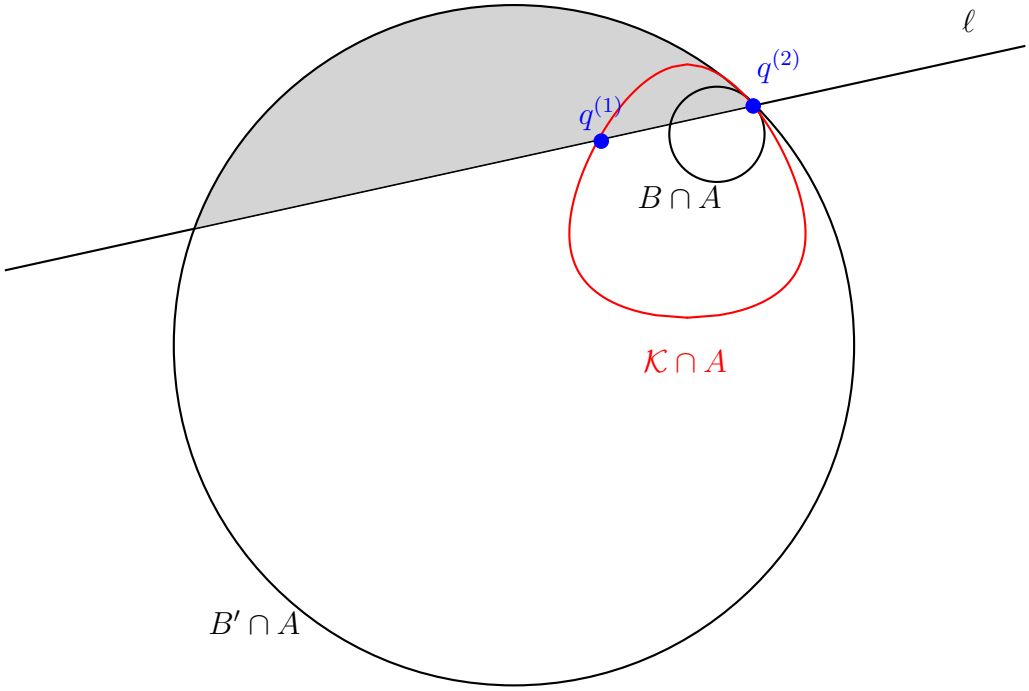

FIGURE 5. The geodesic path from $q^{(1)}$ to $q^{(2)}$ on $\partial\mathcal{K} \cap A$, and the smaller arcs on ∂B and $\partial B'$ connecting the intersections points ℓ with B and B', form the boundary of three nested convex sets $\hat{B} \subseteq \hat{\mathcal{K}} \subseteq \hat{B}'$ (the three nested grey convex regions). Thus, $\text{length}(\partial\hat{B}) \leq \text{length}(\partial\hat{\mathcal{K}}) \leq \text{length}(\partial\hat{B}')$, implying that $\text{dist}_{\partial\mathcal{K}}(\ell \cap \partial\mathcal{K}) \leq \text{dist}_{\partial B'}(\ell \cap \partial B') + \text{length}(\ell \cap B')$.

□

Lemma 6. *Let $\mathcal{K}$ be a convex body and $\mathcal{M} = \partial\mathcal{K}$ be its boundary. Fix $\theta > 0$, and let $\psi^\star_{\theta;(x,v)} := (\gamma^\star_{\theta;(x,v)}, \varphi^\star_{\theta;(x,v)})$ and $\Delta^\star_{\theta;(x,v)}$ be generated by Algorithm 3. Then*

$$\text{dist}(\gamma_{(x,v)}(\Delta^\star_{\theta;(x,v)}), \gamma^\star_{(x,v)}) \leq \alpha \cdot \theta \cdot \Delta^\star_{\theta;(x,v)} \tag{72}$$

and

$$\|\varphi^\star_{(x,v)} - \overline{\varphi_{(x,v)}(\Delta^\star_{\theta;(x,v)})}\| \leq \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta^\star_{\theta;(x,v)} \tag{73}$$

where $\alpha = 2\pi\frac{\mathfrak{M}_2}{\mathfrak{m}_2}$ and $\beta = 5\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}$, as long as $\beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta^\star_{\theta;(x,v)} < 1$.

Moreover, the function $\Delta^\star_{\theta;(x,v)}$ satisfies Assumption 1, with

$$\frac{\theta}{2\sqrt{\mathfrak{M}_2}} \leq \Delta^\star_{\theta;(x,v)} \leq \frac{\theta}{\sqrt{\mathfrak{m}_2}}. \tag{74}$$

Proof. Let ℓ, $\Delta^\star_{\theta;(x,v)}$, θ, and $\psi^\star_{(x,v)}$ be as in Algorithm 3.

(1) Recall that Algorithm 3 sets $\Delta^\star_{\theta;(x,v)} = \text{length}(\ell \cap \mathcal{K})$. Since we can place a tangent sphere to $\partial\mathcal{K}$ at x of radius $\frac{1}{\sqrt{\mathfrak{M}_2}}$ contained in $\mathcal{K}$, and a tangent sphere to $\partial\mathcal{K}$ at x of radius $\frac{1}{\sqrt{\mathfrak{m}_2}}$ containing $\mathcal{K}$ (Figure 5), it must be true that

$$\frac{\theta}{2\sqrt{\mathfrak{M}_2}} \le \frac{1}{\sqrt{\mathfrak{M}_2}}\sin(\theta) \le \Delta^\star_{\theta;(x,v)} \le \frac{1}{\sqrt{\mathfrak{m}_2}}\sin(\theta) \le \frac{\theta}{\sqrt{\mathfrak{m}_2}} \tag{75}$$

Hence, the function $\Delta^\star_{\theta;(x,v)}$ satisfies Assumption 1.

(2) Define $(x_1, v_1) := \psi_{(x,v)}(\Delta^\star_{\theta;(x,v)})$, and $(x_1^\dagger, v_1^\dagger) := \psi^\star_{(x,v)}$. Also define

$$\hat\gamma^\star_{(x,v)}(t) := x + \frac{\gamma^\star_{(x,v)} - x}{\|\gamma^\star_{(x,v)} - x\|} \cdot t \qquad \forall t \ge 0 \tag{76}$$

to be a parametrization of the line segment $\ell \cap \mathcal{K} \subset \mathbb{R}^{d+1}$. In particular, $\hat\gamma^\star_{(x,v)}$ has unit speed, $\hat\gamma^\star_{(x,v)}(0) = x$, and $\hat\gamma^\star_{(x,v)}(\Delta^\star_{\theta;(x,v)}) = \gamma^\star_{(x,v)}$ (since $\Delta^\star_{\theta;(x,v)} = \text{length}(\ell \cap \mathcal{K}) = \|\gamma^\star_{(x,v)} - x\|$).

Since the sectional curvature of the manifold, and hence the curvature of the geodesic trajectory $\gamma_{(x,v)}(t)$, is bounded above by $\sqrt{\mathfrak{M}_2}$, it must be true that

$$\|\hat\gamma^\star_{(x,v)}(t) - \gamma_{(x,v)}(t)\| \le \frac{1 - \cos(\sqrt{\mathfrak{M}_2}t)}{\sqrt{\mathfrak{M}_2}} \le \frac{(\sqrt{\mathfrak{M}_2}t)^2}{\sqrt{\mathfrak{M}_2}} = \sqrt{\mathfrak{M}_2}\cdot t^2 \tag{77}$$

for every $t \ge 0$. Hence by Lemma 5,

$$\text{dist}(\gamma_{(x,v)}(t), \hat\gamma^\star_{(x,v)}(t)) \le 2\pi\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\|\hat\gamma^\star_{(x,v)}(t) - \gamma_{(x,v)}(t)\|$$

$$\le 2\pi\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\sqrt{\mathfrak{M}_2}\cdot t^2 \tag{78}$$

for every $t \ge 0$.

Hence,

$$\text{dist}(\gamma_{(x,v)}(\Delta^\star_{\theta;(x,v)}), \gamma^\star_{(x,v)}) = \text{dist}(\gamma_{(x,v)}(\Delta^\star_{\theta;(x,v)}), \hat\gamma^\star_{(x,v)}(\Delta^\star_{\theta;(x,v)}))$$

$$\le 2\pi\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\sqrt{\mathfrak{M}_2}\cdot(\Delta^\star_{\theta;(x,v)})^2$$

$$\le 2\pi\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\sqrt{\mathfrak{M}_2}\cdot\Delta^\star_{\theta;(x,v)}\cdot\frac{\theta}{\sqrt{\mathfrak{m}_2}}$$

$$= 2\pi\left(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\right)^2 \cdot \Delta^{\star}_{\theta;(x,v)} \cdot \theta$$
$$= \alpha \cdot \Delta^{\star}_{\theta;(x,v)} \cdot \theta$$

(3) Denote by $H_x(v)$ the magnitude of the second derivative of a geodesic with velocity $\frac{v}{\|v\|}$ at $x \in \mathcal{M}$. By Assumption 3, $\sqrt{\mathfrak{m}_2} \leq H_x(v) \leq \sqrt{\mathfrak{M}_2}$. Define $\mathtt{n}(y)$ to be the unit normal vector to $\mathcal{M}$ at y pointing into $\mathcal{K}$, and denote by "$\mathrm{D}_u w(y)$" the directional derivative in the direction u of a vector-valued function $w : \mathcal{M} \mapsto \mathbb{R}^{d+1}$. We note that $\|\mathrm{D}_u \mathtt{n}(y)\| \leq \sqrt{\mathfrak{M}_2}$ at every $y \in \mathcal{M}$ and $u \in \mathbb{S}(\mathcal{T}_y)$. Denote by $\mathcal{P}$ and $\mathcal{P}^{\perp}$ the operators on $\mathbb{R}^{d+1}$ that, respectively, project a vector onto the 2-plane P, and the orthogonal complement $P^{\perp}$ of the 2-plane P, where P is the plane spanned by the vector $\mathtt{n}(x)$ and the line ℓ (this 2-plane is also the 2-plane of curvature associated with the curve $\gamma_{(x,v)}(t)$ at $t = 0$).

Define $(x(t), v(t)) := \psi_{(x,v)}(t)$. Since $\mathcal{P}^{\perp} v(0) = 0$ and $\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(0)) = \mathcal{P}^{\perp} v = 0$, we have:

$$\|\mathcal{P}^{\perp} v(t)\| = \|\mathcal{P}^{\perp} v(t) - \mathcal{P}^{\perp} v(0)\| = \left\| \int_0^t \mathcal{P}^{\perp} H_{x(\tau)}(v(\tau))\mathtt{n}(\gamma_{(x,v)}(\tau))\mathrm{d}\tau \right\|$$
$$\leq \int_0^t \left\| \mathcal{P}^{\perp} H_{x(\tau)}(v(\tau))\mathtt{n}(\gamma_{(x,v)}(\tau)) \right\| \mathrm{d}\tau$$
$$= \int_0^t \left\| H_{x(\tau)}(v(\tau))\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(\tau)) \right\| \mathrm{d}\tau$$
$$\leq \int_0^t \left\| \sqrt{\mathfrak{M}_2}\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(\tau)) \right\| \mathrm{d}\tau$$
$$= \sqrt{\mathfrak{M}_2} \int_0^t \left\| \int_0^\tau \mathrm{D}_{v(s)}\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(s))\mathrm{d}s + \mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(0)) \right\| \mathrm{d}\tau$$
$$= \sqrt{\mathfrak{M}_2} \int_0^t \left\| \int_0^\tau \mathrm{D}_{v(s)}\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(s))\mathrm{d}s + 0 \right\| \mathrm{d}\tau$$
$$\leq \sqrt{\mathfrak{M}_2} \int_0^t \int_0^\tau \left\| \mathrm{D}_{v(s)}\mathcal{P}^{\perp}\mathtt{n}(\gamma_{(x,v)}(s)) \right\| \mathrm{d}s\mathrm{d}\tau$$
$$\leq \sqrt{\mathfrak{M}_2} \int_0^t \int_0^\tau \sqrt{\mathfrak{M}_2}\mathrm{d}s\mathrm{d}\tau$$
$$= \frac{(\sqrt{\mathfrak{M}_2})^2}{2} t^2. \tag{79}$$

Therefore, since $\mathcal{P}^\perp v_1^\dagger = 0$, Equation 79 gives

(80)
$$\|\mathcal{P}^\perp v_1^\dagger - \mathcal{P}^\perp v_1\| = \|\mathcal{P}^\perp v_1\| = \|\mathcal{P}^\perp v(\Delta^\star_{\theta;(x,v)})\| \leq \frac{(\sqrt{\mathfrak{M}_2})^2}{2}(\Delta^\star_{\theta;(x,v)})^2.$$

Also, observe that the change in a vector's orientation in ambient Euclidean space when undergoing parallel transport along a geodesic path is maximized when the vector points in the same direction as that geodesic path. Therefore, by the same calculation as in Equation 79,

(81)
$$\|\mathcal{P}^\perp v_1 - \mathcal{P}^\perp \overline{v}_1\| \leq \left\| \int_0^{\bar{t}} \mathcal{P}^\perp H_{x(\tau)}(v(\tau))\mathfrak{m}(\gamma_{(x,v)}(\tau))\mathrm{d}\tau \right\| \leq \frac{(\sqrt{\mathfrak{M}_2})^2}{2}\bar{t}^2$$

for all $t > 0$, where $\overline{v}_1$ can be defined to be the parallel transport from of v_1 from x_1 to $x_1^\dagger$ along *any* distance minimizing geodesic, and $\bar{t} := \mathrm{dist}(x_1, x_1^\dagger)$. Since $\gamma_{(x,v)}(t)$ has unit speed and $\hat{\gamma}^\star_{(x,v)}(t)$ has speed at most 1, it must be true that $\bar{t} \leq 2\Delta^\star_{\theta;(x,v)}$. Therefore,

(82)
$$\|\mathcal{P}^\perp v_1 - \mathcal{P}^\perp \overline{v_1}\| \leq \frac{(\sqrt{\mathfrak{M}_2})^2}{2} \cdot 4(\Delta^\star_{\theta;(x,v)})^2.$$

Hence, Equations 80 and 82 imply that

$$\|\mathcal{P}^\perp v_1^\dagger - \mathcal{P}^\perp \overline{v_1}\| \leq \left[\|\mathcal{P}^\perp v_1^\dagger - \mathcal{P}^\perp v_1\| + \|\mathcal{P}^\perp v_1 - \mathcal{P}^\perp \overline{v_1}\|\right]$$

$$\leq \left[\frac{(\sqrt{\mathfrak{M}_2})^2}{2} \cdot (\Delta^\star_{\theta;(x,v)})^2 + \frac{(\sqrt{\mathfrak{M}_2})^2}{2} \cdot 4(\Delta^\star_{\theta;(x,v)})^2\right]$$

(83)
$$= \frac{5}{2}(\sqrt{\mathfrak{M}_2})^2 \cdot (\Delta^\star_{\theta;(x,v)})^2$$

for every $\theta > 0$, and for every choice of distance-minimizing geodesic for which one can define the parallel transport "$\overline{v_1}$".

In this paragraph, we will write $v_1^{\dagger\theta} \equiv v_1^\dagger$ and $\overline{v_1}^\theta \equiv \overline{v_1}$ to specify the value of θ. First, observe that, for every θ, $v_1^{\dagger\theta}$ and $\mathcal{P}\overline{v_1}^\theta$ both lie in the 2-plane P, and are both tangent to $\mathcal{M}$, so they both must be tangent to the 1-dimensional curve $P \cap \mathcal{M}$ at the same point $x_1^{\dagger\theta}$. Therefore, $v_1^{\dagger\theta}$ and $\mathcal{P}\overline{v_1}^\theta$ both lie on the same line. Since this holds for every θ, and every choice of distance-minimizing geodesic, $v_1^{\dagger\theta}$ and $\overline{v_1}^\theta$ must be pointing in the same direction, provided that $5(\sqrt{\mathfrak{M}_2})^2 \cdot (\Delta^\star_{\tilde{\theta};(x,v)})^2 < 1$ for every $\tilde{\theta} \leq \theta$, since we can choose a family of geodesics that is continuous in $\tilde{\theta}$ to define the $\overline{v_1}^{\tilde{\theta}}$ at every $\tilde{\theta} \leq \theta$ (since then

$\|\mathcal{P}^{\perp}v_1^{\dagger\tilde{\theta}} - \mathcal{P}^{\perp}\overline{v_1}^{\tilde{\theta}}\| < 1$ for all $\tilde{\theta} \leq \theta$, and hence $\|v_1^{\dagger\tilde{\theta}} - \mathcal{P}\overline{v_1}^{\tilde{\theta}}\| \neq 0$ for all $\tilde{\theta} \leq \theta$. By continuity, $v_1^{\dagger\tilde{\theta}}$ and $\mathcal{P}\overline{v_1}^{\tilde{\theta}}$ must be pointing in the same direction for all $\tilde{\theta} \leq \theta$, since they point in the same direction when $\tilde{\theta} = 0$). Therefore $v_1^{\dagger}$ and $\overline{v_1}$ must be pointing in the same direction.

Therefore, since $\overline{v_1}$ and $v_1^{\dagger}$ both have unit magnitude, and the projection $\mathcal{P}\overline{v_1}$ has the same direction as $v_1^{\dagger}$, we must have that

$$\|v_1^{\dagger} - \overline{v_1}\| \leq 2\|\mathcal{P}^{\perp}\overline{v_1}\|$$

$$= 2\|\mathcal{P}^{\perp}v_1^{\dagger} - \mathcal{P}^{\perp}\overline{v_1}\|$$

$$\leq 5(\sqrt{\mathfrak{M}_2})^2 \cdot (\Delta^{\star}_{\theta;(x,v)})^2. \tag{84}$$

Thus,

$$\|\varphi^{\star}_{(x,v)} - \overline{\varphi_{(x,v)}(\Delta^{\star}_{\theta;(x,v)})}\| \leq 5(\sqrt{\mathfrak{M}_2})^2 \cdot (\Delta^{\star}_{\theta;(x,v)})^2 \tag{85}$$

$$\leq 5(\sqrt{\mathfrak{M}_2})^2 \cdot \frac{\theta}{\sqrt{\mathfrak{m}_2}} \cdot \Delta^{\star}_{\theta;(x,v)}$$

$$= \beta \cdot \sqrt{\mathfrak{M}_2} \cdot \theta \cdot \Delta^{\star}_{\theta;(x,v)}$$

□

We conclude with a bound on the total computational effort that is required to sample approximately from the uniform distribution on $\partial\mathcal{K}$ with error less than any fixed $\epsilon > 0$:

Theorem 7. *Fix $\epsilon > 0$, let α, β be defined as in Lemma 6, and let $I(\epsilon), \theta(\epsilon)$ be defined as in Equation 59. Let $\{X_i^{\theta(\epsilon)}\}_{i\in\mathbb{N}}$ be the Markov chain defined in Equation 41, with the oracle ★ provided by Algorithm 3. Then the distribution $\Pi^{\theta(\epsilon)}_{I(\epsilon)+1}$ of the $X^{\theta(\epsilon)}_{I(\epsilon)}$ satisfies*

$$W_d(\Pi^{\theta(\epsilon)}_{I(\epsilon)+1}, \Pi) \leq \epsilon \tag{86}$$

and the intersection oracle and normal vector oracles in Algorithm 3 are called no more than

$$N(\epsilon) = \mathcal{O}\left((\frac{\mathfrak{M}_2}{\mathfrak{m}_2})^3 \times \frac{\log(\frac{2\pi}{\epsilon\sqrt{\mathfrak{m}_2}})}{\epsilon \cdot \sqrt{\mathfrak{M}_2}}\right) \tag{87}$$

times.

Proof. Inequality 86 is an application of Lemma 4. To prove Inequality 87, by Theorem 6,

$$
\begin{aligned}
N(\epsilon) &\leq \frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \left[\lceil \frac{\pi\sqrt{\mathfrak{M}_2}}{2\sqrt{\mathfrak{M}_2}\,\theta(\epsilon)}\rceil + \mathcal{O}(\log(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}))\right] \\
&= \mathcal{O}\Bigg(\frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \\
&\quad \times \left[\frac{1}{\epsilon\cdot\sqrt{\mathfrak{M}_2}(1-\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \left[1+\frac{\pi}{2}\alpha+(\frac{\pi}{2})^2\beta\right] + \mathcal{O}(\log(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}))\right]\Bigg) \\
&= \mathcal{O}\Bigg(\frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))}\times \\
&\quad \left[\frac{1}{\epsilon\cdot\sqrt{\mathfrak{M}_2}(1-\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \left[1+\frac{\pi}{2}\cdot 2\pi(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}})^2+(\frac{\pi}{2})^2\cdot\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}\right] + \mathcal{O}(\log(\frac{\mathfrak{M}_2}{\mathfrak{m}_2}))\right]\Bigg) \\
&= \mathcal{O}\left(\frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\log(\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))} \times \frac{1+\frac{\pi}{2}\cdot 2\pi(\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}})^2+(\frac{\pi}{2})^2\cdot\frac{\sqrt{\mathfrak{M}_2}}{\sqrt{\mathfrak{m}_2}}}{\epsilon\cdot\sqrt{\mathfrak{M}_2}(1-\cos(\frac{\pi\sqrt{\mathfrak{m}_2}}{2\sqrt{\mathfrak{M}_2}}))}\right) \\
&= \mathcal{O}\left((\frac{\mathfrak{M}_2}{\mathfrak{m}_2})^3 \times \frac{\log(\frac{\epsilon\sqrt{\mathfrak{m}_2}}{2\pi})}{\epsilon\cdot\sqrt{\mathfrak{M}_2}}\right),
\end{aligned}
$$

where the first inequality is due to Inequality 63 and Lemma 6, and the remaining lines are obtained by the definitions of $\theta(\epsilon)$, α and β. □

Remark 2. *If the smoothness of the curvature of $\mathcal{M}$ is uniformly bounded below (i.e., if the derivative of the curvature is uniformly bounded above), we can replace the ratio $\frac{\mathfrak{M}_2}{\mathfrak{m}_2}$ appearing in Lemmas 5 and 6 by a constant $\mathfrak{M}_3$ that depends only on a lower bound on the smoothness of the curvature, not on the curvature itself, so that we get $\alpha = 2\pi\mathfrak{M}_3$ and $\beta = \sqrt{\mathfrak{M}_3}$. The number of intersection and normal line oracle calls is then $\mathcal{O}^*((\frac{\mathfrak{M}_2}{\mathfrak{m}_2})^2\cdot\mathfrak{M}_3\cdot(\epsilon\sqrt{\mathfrak{M}_2})^{-1})$, which is quadratic in the curvature ratio $\frac{\mathfrak{M}_2}{\mathfrak{m}_2}$.*

9. Acknowledgements

We are grateful to Yin Tat Lee and A.B. Dieker for helpful discussions. Oren Mangoubi was supported by a Canadian Statistical Sciences Institute (CANSSI) Postdoctoral Fellowship, and by an NSERC Discovery grant. Aaron Smith was supported by an NSERC Discovery grant.

References

[1] Uwe Abresch and Wolfgang T. Meyer. Injectivity radius estimates and sphere theorems. *Comparison geometry*, 30:1–47, 1997.

[2] Claude JP Bélisle, H Edwin Romeijn, and Robert L Smith. Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18(2):255–266, 1993.

[3] Francis Bonahon and Yasar Sözen. Variation of the Liouville measure of a hyperbolic surface. *Ergodic Theory and Dynamical Systems*, 23(03):729–758, 2003.

[4] Jeff Cheeger and David G. Ebin. *Comparison theorems in Riemannian geometry*, volume 365. American Mathematical Soc., 2008.

[5] Francis Comets, Serguei Popov, Gunter M Schütz, and Marina Vachkovskaia. Billiards in a general domain with random reflections. *Archive for rational mechanics and analysis*, 191(3):497–537, 2009.

[6] Persi Diaconis, Susan Holmes, and Mehrdad Shahshahani. Sampling from a manifold. In *Advances in Modern Statistical Theory and Applications: A Festschrift in Honor of Morris L. Eaton*, pages 102–125. Institute of Mathematical Statistics, 2013.

[7] A.B. Dieker and Santosh S. Vempala. Stochastic billiards for sampling from the boundary of a convex set. *Mathematics of Operations Research*, 40(4):888–901, 2015.

[8] Jost-Hinrich Eschenburg. *Comparison theorems in Riemannian geometry*. Università delli studi di Trento, Dipartimento di matematica, 1994.

[9] Sigmundur Gudmundsson. *An Introduction to Riemannian Geometry*. Lund University, 2008.

[10] Heinz Hopf and Willi Rinow. Über den begriff der vollständigen differentialgeometrischen fläche. *Commentarii Mathematici Helvetici*, 3(1):209–225, 1931.

[11] Wilhelm PA Klingenberg. *Riemannian geometry*, volume 1. Walter de Gruyter, 1995.

[12] Gilles Lebeau and Laurent Michel. Semi-classical analysis of a random walk on a manifold. *The Annals of Probability*, 38(1):277–315, 2010.

[13] Sumner B Myers. Riemannian manifolds with positive mean curvature. *Duke Math. J*, 8(40):1–404, 1941.

[14] Yann Ollivier. Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis*, 256(3):810–864, 2009.

[15] A. Cengiz Öztireli, Marc Alexa, and Markus Gross. Spectral sampling of manifolds. *ACM Transactions on Graphics (TOG)*, 29(6):168, 2010.
[16] Peter Petersen. *Riemannian geometry*, volume 171. Springer, 2006.
[17] Harry Ernest Rauch. A contribution to differential geometry in the large. *Annals of Mathematics*, pages 38–55, 1951.
[18] Chang-han Rhee, Enlu Zhou, and Peng Qiu. An iterative algorithm for sampling from manifolds. In *Proceedings of the 2014 Winter Simulation Conference*, pages 574–585. IEEE Press, 2014.
[19] Yin Tat Lee and Santosh S Vempala. Geodesic walks on polytopes. *arXiv preprint arXiv:1606.04696*, 2016.
[20] Santosh Vempala. Geometric random walks: a survey. *Combinatorial and computational geometry*, 52(573-612):2, 2005.
[21] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
[22] Paweł G. Walczak. Existence of smooth invariant measures for geodesic flows of foliations of riemannian manifolds. *Proceedings of the American Mathematical Society*, 120(3):903–906, 1994.